FRONT /› CODE

NEXT-GENERATION PHISHING

APRIL 2026

www.front-code.com

Next-Generation Phishing in Cybersecurity

Identity, Phishing, Social Engineering, Credential Harvesting

April, 2026

Audience: Security-aware leaders, engineers, and students

Scope: Defensive, non-operational guidance (no offensive instructions)

FRONT-CODE.COM

Next-Generation Phishing in Cybersecurity 2019–2026

Executive summary

Phishing from 2019–2026 is best understood as a shift from "fake login pages + stolen passwords" toward **identity-flow abuse, token theft, and interactive social engineering** across channels (email, SMS, collaboration tools, QR codes, and voice/video). Threat actors increasingly avoid classic indicators (obvious typos, suspicious domains, credential form posts) and instead **weaponize legitimate authentication infrastructure** (OAuth/SSO flows, device authorization grants, error redirects, and SaaS "connected apps") so that the victim's actions produce valid tokens or trusted changes. Microsoft's telemetry highlights layered campaigns combining **device code phishing** and **OAuth consent phishing**, sometimes redirecting victims into **adversary-in-the-middle (AiTM)** sites, and notes rapid adoption patterns for device code phishing. [1]

Across incident response observations, defenders are also seeing a **rebalancing of social engineering modalities**: Mandiant reports that as automated technical controls improved, **email phishing fell to 6% of intrusions (in Mandiant's 2025 dataset)** and adversaries pivoted toward more interactive methods, including voice-based social engineering. [2] In parallel, Google's cloud-oriented reporting highlights meaningful shares of cases involving **vishing**, including helpdesk impersonation and SaaS identity takeover patterns. [3]

Three technical evolutions materially change detection and response:

1. **Session and token theft over password theft.** AiTM phishing kits (reverse proxies) capture credentials *and* MFA outputs and steal session cookies, enabling logins that look legitimate in many logs. Microsoft describes Tycoon2FA as an AiTM phishing kit operated at scale, lowering the barrier for MFA bypass; it was disrupted through a multi-party operation announced March 4, 2026. [4]
2. **OAuth abuse beyond "fake login pages."** Device code phishing abuses the OAuth device authorization grant (RFC 8628) and can provide access without collecting a password, while newer patterns exploit OAuth redirect behavior to route victims from legitimate login URLs to attacker-controlled infrastructure without token theft. [5]
3. **Out-of-band and cross-device delivery.** QR-code phishing ("quishing") and mobile-driven workflows bypass some email link controls; PDF attachments with QR codes (e.g., "Scanception") and QR shortener ecosystems are repeatedly noted in threat reporting. [6]

Assumptions (explicit): - Target audience: cybersecurity professionals (SOC analysts, security engineers, incident responders, risk/GRC leaders).
- Scope: global (public + private sector), with examples drawn from widely reported cases and vendor/government reporting.
- Timeframe: January 1, 2019 through April 7, 2026 (current date, Europe/Berlin).
- Definition of phishing: inclusive of social-engineering attacks intended to induce unsafe actions or disclosure (credentials, tokens, approvals, payments), including BEC and voice/video deception (consistent with industry reporting). [7]

- “Detection efficacy” in tables is qualitative (High/Medium/Low) because real-world efficacy is environment- and implementation-dependent; metrics vary across vendors and datasets.

Taxonomy of modern phishing methods 2019–2026

A useful taxonomy needs to classify phishing by (a) **delivery channel**, (b) **what the attacker is trying to obtain** (credentials, tokens, approvals, money, execution), and (c) **where trust is subverted** (email identity, web identity, human authority, process compliance). Industry trend reporting underscores how large-scale phishing continues even as individual lures evolve; the Anti-Phishing Working Group[8] reports 3.8 million phishing attacks observed in 2025 (slightly above 2024) and sustained high quarterly volumes. [9]

Comparative taxonomy table

Modern phishing method	Primary channels	Primary objective	“New” element (2019–2026 emphasis)	Typical downstream impact
AI-generated phishing content (LLM-assisted)	Email, SMS, chat/collab, social DM	Increase conversion via personalization	Role- and context-aligned pretexts at scale; improved language quality reduces “typo detection” value; enables rapid A/B iteration	Credential/token theft, malware delivery, BEC enablement
Deepfake-enabled impersonation	Voice calls, video calls, voicemail, media messages	Authority hijack (approvals, payments, recovery)	Voice cloning/video deepfakes increase plausibility of “executive/helpdesk” pretexts	Fraudulent transfers, account recovery bypass, helpdesk resets
Vishing (voice phishing) & helpdesk social engineering	Phone, voicemail, callback scams	Credential reset / MFA reset / app authorization / payment	Interactive manipulation; targets service desks and SaaS admins; often paired with “visit this URL now”	SaaS takeover, data theft + extortion, ransomware staging
Smishing (SMS phishing)	SMS, messaging apps	Credential theft; MFA capture; link clicks	Mobile-first targeting; corporate phones often outside endpoint controls; short links	Account takeover, lateral movement via SSO
OAuth device code phishing	Email, chat, SMS, voice-assisted	Token theft without password capture	Abuses legitimate OAuth device flow (RFC 8628); victim enters code on real	Persistent account access, data theft, internal phishing

Modern phishing method	Primary channels	Primary objective	"New" element (2019–2026 emphasis)	Typical downstream impact
			login page; attacker receives tokens	
OAuth consent / connected-app phishing	Email, web, voice-assisted	Victim grants OAuth scopes	"By-design" permission grants bypass MFA and can persist beyond password resets	Long-lived access; mailbox rule abuse; SaaS data exfiltration
OAuth redirection abuse	Email links	Drive victim to attacker infrastructure via "legit" login URL	Uses silent OAuth flows + invalid scopes to trigger redirects without stealing tokens	Malware delivery, follow-on compromise
AiTM / reverse-proxy phishing (MFA bypass kits)	Email, web	Session cookie theft + MFA interception	Proxy captures credentials and session tokens in real time; commoditized "phishing-as-a-service"	High-confidence account takeover (ATO), mailbox manipulation, BEC
QR-code phishing ("quishing")	Email attachments (PDF/images), posters, invoices	Offload click to mobile scan	QR codes evade some link scanners; cross-device context switching reduces user scrutiny	Credential theft, malware, MFA capture
Homograph/IDN lookalike domains	Email, web, ads	Brand impersonation	Unicode confusables and IDNA complexity; visual similarity defeats casual inspection	Credential theft, malware, invoice fraud
BEC / vendor email compromise / invoice fraud	Email, sometimes voice follow-up	Money movement	Increasingly pairs mailbox takeover + process manipulation; may leverage deepfakes for confirmation	Direct financial loss; regulatory/reporting exposure
Supply-chain phishing	Email, shared SaaS, vendor platforms	Compromise via trusted third party	Trusted sender compromise (vendor mailbox, marketing platform, SaaS integration)	Multi-org spread, invoice fraud, initial access brokerage
Multi-stage / polymorphic campaigns	Any/all	Resilience against controls	Rotating kits, dynamic redirects, CAPTCHAs/"human	Blended outcomes: ATO → BEC → extortion/ransomware

Modern phishing method	Primary channels	Primary objective	“New” element (2019–2026 emphasis)	Typical downstream impact
			verification,” staged payloads	

This table emphasizes that “modern phishing” is frequently **identity-centric**: the objective is to obtain **identity artifacts** (tokens, sessions, app grants) or to induce **trusted actions** (wire transfers, MFA resets), rather than merely to collect a password.

Technical mechanisms and attack chains

Modern phishing attack chains commonly follow an identifiable sequence: **targeting → pretext delivery → trust/identity subversion → privilege expansion → monetization**. The novelty is in the identity subversion step: attackers increasingly co-opt legitimate identity flows so the “malicious” part is the victim’s decision, not a suspicious login form.

Email authentication and why it’s insufficient alone

Email authentication (SPF, DKIM, DMARC) hardens a domain’s ability to assert what infrastructure is authorized to send on its behalf. National Institute of Standards and Technology[10] explicitly describes SPF, DKIM, and DMARC as mechanisms that help receivers evaluate authenticity and sender policy. [11] The standards define the mechanics: SPF (RFC 7208) authorizes sending hosts, DKIM (RFC 6376) signs messages, and DMARC (RFC 7489) expresses domain policy and alignment expectations. [12]

However, these controls do **not** prevent: - Phishing from **lookalike domains** (close spelling),
- Phishing from **compromised legitimate accounts** (passes authentication),
- Many **social engineering** attacks that don’t rely on spoofing, and
- Cross-channel flows (QR, voice) that bypass email-only checks. [13]

AiTM phishing: reverse proxies and “session as the prize”

AiTM phishing inserts a reverse proxy between the victim and the legitimate login service. The victim sees a believable page; the proxy relays credentials and MFA to the real service and steals the resulting session artifacts (cookies/tokens). Microsoft documents a large-scale AiTM kit ecosystem in its coverage of Tycoon2FA and notes these kits can allow less skilled actors to bypass MFA. [14]

A key operational point for defenders: AiTM-driven compromise can produce authentication events that appear “normal” (correct MFA, expected geolocation if the attacker replays tokens appropriately), increasing reliance on **post-authentication telemetry** (mailbox rule changes, OAuth grants, anomalous access to files, new forwarding addresses).

OAuth device code phishing: abusing RFC 8628 for token capture

RFC 8628 defines the OAuth 2.0 device authorization grant intended for input-constrained devices (e.g., smart TVs). [15] Attackers weaponize this by generating a device code, sending it to a victim, and persuading the victim to enter it at the legitimate device-login page. The victim’s successful authentication yields tokens that the attacker can capture and use.

Microsoft’s 2025 reporting describes device code phishing as high risk precisely because it enables access **without a password** and can capture access/refresh tokens, and also notes rapid adoption (with

a large share observed in the latter half of the observation period). [16] Microsoft further reports operationalization into collaboration contexts (e.g., presenting device codes inside Teams-like lures), which erodes users' ability to identify "something unusual." [17]

Proofpoint reports criminal adoption as well, observing high-volume credential phishing actors abusing OAuth device code authorization for account takeover beginning in October 2025, illustrating diffusion from state-aligned use into financially motivated operations. [18]

In April 2026, Microsoft described an **AI-enabled device code phishing campaign** that used generative AI for role-aligned lures and automated device code generation "just in time" to evade expiration windows (15 minutes) by triggering code creation when the user clicked. [19] This pattern is a concrete example of how automation targets the *mechanics* of identity flows, not just message wording.

OAuth consent and connected-app phishing: "by-design" persistence

OAuth consent phishing persuades a user to grant an attacker-controlled application permissions (scopes). Microsoft describes app consent phishing as bypassing MFA and persisting beyond password resets, and notes attackers pivoting to "workload identities" (apps/services/scripts) as phishing-resistant MFA and conditional access improve. [20]

A parallel pattern appears in SaaS compromises: the Google Threat Intelligence Group describes a financially motivated cluster (UNC6040) that uses vishing to persuade victims to authorize a malicious connected app to a Salesforce portal (often a modified Data Loader), enabling data theft and later extortion. [21]

OAuth redirection abuse: "legitimate login URL → malicious destination"

In March 2026, Microsoft described phishing-led exploitation of OAuth redirection mechanisms using silent OAuth flows and intentionally invalid scopes to redirect victims to attacker infrastructure **without stealing tokens** in that flow. [22] This matters because many users and some controls implicitly trust a "legitimate login domain" URL; the attacker leverages that trust to get browser execution or malware download downstream.

QR-code phishing and PDF attachment campaigns

The Anti-Phishing Working Group[8] explicitly highlighted "Millions of Malicious QR Codes Lead to Phishing" in its Q1 2025 reporting, noting that QR codes in emails lead to phishing sites and malware. [23]

Campaign reporting illustrates why QR works operationally: - The email body may have no clickable URL (reducing URL rewriting/link scanning triggers).
- The "click" occurs on a phone camera or QR scanner, potentially outside enterprise protections.
- QR codes embedded in PDFs can evade simplistic content checks.

Broadcom's description of "Scanception" frames it as QR codes embedded in PDF attachments that redirect users to fake Microsoft 365 login pages for credential harvesting. [24] Palo Alto Networks Unit 42 reports increasing QR code shortener traffic and quantifies growth across half-year windows, indicating an expanding infrastructure ecosystem around QR-driven redirection. [25]

ClickFix / “copy-paste execution” as a phishing-adjacent chain

ClickFix attacks are a modern example of social engineering that aims not for credentials but for **user-driven code execution**. Microsoft’s August 21, 2025 analysis describes ClickFix as leading victims to a lure page and tricking them into executing a malicious command themselves, helping bypass conventional automated security solutions by adding a human-in-the-loop step. [26] MITRE ATT&CK codifies this pattern as “User Execution: Malicious Copy and Paste.” [27]

Mermaid flowchart: a composite modern chain (AI-enabled device code phishing → token theft → BEC)

```
flowchart TD
  A[OSINT + role context collection] --> B[LLM-assisted lure generation<br/>RFP/invoice/workflow theme]
  B --> C[Delivery: email/chat/SMS with "Verify device" pretext]
  C --> D[Victim clicks link]
  D --> E[Attacker triggers just-in-time device code generation]
  E --> F[Victim enters code on legitimate device login page]
  F --> G[Legitimate IdP issues access/refresh tokens]
  G --> H[Attacker captures tokens and establishes session]
  H --> I[Post-auth actions: mailbox search + exfil]
  I --> J[Persistence: create inbox rules / add OAuth app consent]
  J --> K[Internal phishing + vendor impersonation]
  K --> L[BEC: invoice/wire change request]
```

Key elements in this composite chain are grounded in reported behavior: AI-assisted role-aligned lures and just-in-time device code generation (Microsoft, April 6, 2026), device code phishing’s token-based access (Microsoft reporting), and follow-on actions typical of ATO→BEC playbooks. [28]

Social-engineering vectors and psychological triggers

Technical sophistication does not replace social engineering; it **repackages** it around workflows users already trust (SSO prompts, helpdesk calls, compliance tasks). Recent research increasingly models phishing through the lens of **cognitive biases and persuasion principles**, not just message features.

Core psychological levers (with research grounding)

Open-access research on cognitive bias exploitation in phishing emails (2025) finds attackers systematically leverage biases and that incorporating cognitive-bias features can improve detection models. [29] A 2024 survey of persuasion principles in phishing similarly emphasizes Cialdini-like constructs (authority, scarcity, reciprocity, social proof, liking, consistency) as recurring manipulation mechanisms. [30]

Operationally, modern campaigns often blend these levers:

- **Authority + legitimacy transfer:** “IT support,” “finance leadership,” or “compliance” roles. Vishing campaigns explicitly impersonate IT/helpdesk staff; Google’s UNC6040 case describes operators impersonating IT support to get SaaS authorization and credentials. [21]
- **Urgency/time pressure:** Short windows, expiring links, payroll deadlines, “security incident in progress.” Research on persuasion under pressure shows time constraints modulate how well users differentiate phishing from genuine emails. [31]

- **Scarcity and fear:** “Account will be locked,” “invoice overdue,” “legal escalation.”
- **Commitment/consistency:** Attackers induce small initial compliance (join a call, scan a QR, enter a code) then escalate (“now approve MFA,” “now install tool,” “now process transfer”).
- **Social proof / familiarity:** “A colleague shared a file,” “Teams invite,” “DocuSign,” “OneDrive,” etc.—themes Microsoft reports in device code phishing lures. [32]
- **Curiosity and habituation:** QR codes in PDFs and “document shared” lures exploit habitual document-handling routines; APWG and vendor reporting emphasize QR as a large-scale lure technique. [33]

Why AI changes the psychology, not just the grammar

The most defensible claim about AI in phishing is not “AI writes better English,” but that AI enables **faster personalization loops** and **multi-channel orchestration**. Microsoft’s April 2026 campaign write-up describes role-aligned themes and automation that adapt timing to device code validity constraints. [19] Research syntheses on LLM-generated phishing campaigns emphasize that LLMs enable stylistically adaptive generation and multilingual reach, undermining detection methods that rely on surface-level linguistic errors. [34]

Deepfakes and vishing: collapsing “out-of-band verification”

Voice and video historically served as a verification channel (“call them to confirm”). Deepfakes and sophisticated vishing erode that assumption. Reuters reported on a 2024 case cited by Hong Kong police where a finance employee transferred over $25 million after joining a video call populated by deepfake “colleagues.” [35] Earlier cases already indicated voice cloning fraud feasibility (e.g., 2019 CEO-voice scams reported in business press). [36]

From a defensive research standpoint, audio deepfake detection faces realism challenges: one 2025 paper argues models often fail to generalize because lab datasets don’t reflect the “presentation channel” (phone, compression, playback), and proposes more realistic evaluation frameworks. [37]

Detection techniques and limitations

Detection needs to be framed as **layered control surfaces**, because modern phishing often “wins” by moving the decisive user action outside a single telemetry domain (email → mobile scan → SaaS consent → token replay).

Email filtering and its blind spots

Secure email gateways, reputation systems, and content classifiers remain important, but modern approaches exploit gaps: - Legitimate login URLs used as part of malicious redirect chains (OAuth redirect abuse). [22]
- Emails without clickable links (QR in attachments) reduce efficacy of URL rewriting/scan-at-click controls. [38]
- Out-of-band social engineering (phone calls) bypasses email controls entirely; Mandiant reports a pivot toward voice-based social engineering in intrusions. [39]

DMARC/SPF/DKIM: strengths and limits (practical view)

Email authentication is still foundational for reducing direct domain spoofing. NIST SP 800-177 provides an architectural view of how receivers compare “From” addresses with SPF/DKIM results and DMARC policy to determine handling and reporting. [11] The underlying standards specify the domain-level nature of these assertions. [12]

But **domain authentication is not sender intent authentication**. Modern campaigns commonly: - Use attacker-controlled domains that look similar (typosquatting/homographs),
- Use compromised legitimate accounts,
- Use SaaS platforms where messages legitimately originate and authenticate, and
- Shift delivery to QR/voice where DMARC is irrelevant. [40]

Behavioral analytics and identity telemetry (ITDR/UEBA style)

Because tokens and sessions can create “legitimate-looking” access, detection must include: - anomalous OAuth app grants (new consent, unusual scopes), - device code sign-in events, - mailbox rule creation/forwarding, - impossible travel / atypical IPs (with caution for VPN), - unusual SaaS data export patterns and bulk downloads.

Microsoft’s OAuth redirect abuse report explicitly notes Microsoft Defender correlating malicious activity across email, identity, and endpoint signals—an example of why cross-domain correlation matters. [22]

For AWS/Azure/M365-like environments, Cybersecurity and Infrastructure Security Agency[41] advisory AA21-008A documents tooling and methods to detect post-compromise activity in Microsoft cloud environments; CISA’s Sparrow tool is designed to detect compromised accounts and applications endemic to identity/authentication-based attacks. [42]

URL and domain analysis, including IDN/homograph detection

URL analysis remains relevant, particularly for: - newly registered domains, - suspicious hosting patterns, - redirect chains, - known phishing kit fingerprints.

Homograph/IDN attacks require additional defenses because Unicode confusables can make malicious domains visually indistinguishable. IDNA standards formalize internationalized domain handling. [43] Unicode security reporting discusses spoofing through visually confusable strings. [44] Academic work continues to develop detection methods; for example, a 2024 paper proposes Siamese neural networks for detecting camouflaged IDN-based phishing and reports high detection accuracy in its evaluated datasets. [45]

QR-code phishing detection: where controls should attach

QR-code phishing creates a key telemetry gap: the email system may never see a URL click. Threat reporting notes substantial QR-code shortener ecosystem growth and QR-in-PDF campaigns. [46] Defenses generally need to attach at: - attachment analysis (extract QR content), - mobile threat defense / managed browser, - DNS/HTTP filtering on mobile, - phishing-resistant authentication so credential reuse is less valuable.

Voice/deepfake detection: emerging but not turnkey

Academic proposals exist for vishing detection (including LLM-augmented audio classification approaches), but data realism and deployment constraints remain substantial. [47] For defenders, the most reliable short-term mitigation is procedural: **out-of-band verification must become "out-of-workflow" verification** (independent callback numbers, verified directories, transaction signing, and step-up authentication), not merely "call them back" if the phone channel itself can be manipulated.

Detection efficacy matrix (qualitative)

Interpretation: "High" means the control class is generally effective *when well-implemented*; "Low" means attackers typically bypass it by design.

Method (from taxonomy)	DMARC/SPF/DKIM	Secure email gateway & URL scanning	Identity telemetry (sign-ins, consent, token use)	Endpoint security (EDR)	User training/process controls	Notes
Classic credential phishing	Medium	High	Medium	Medium	Medium	Still common; controls reasonably mature
AiTM / reverse proxy kits	Low	Medium	High	Medium	Medium	Identity-side anomaly + post-login behavior are decisive [14]
OAuth device code phishing	Low	Medium	High	Low–Medium	Medium	Often no fake login page; requires device-flow monitoring [48]
OAuth consent / connected-app phishing	Low	Medium	High	Low	Medium	App governance and consent review are central [49]
OAuth redirection abuse	Low	Medium	Medium	High	Medium	Endpoint + web controls matter because destination is malicious [50]
QR-code phishing	Low	Low–Medium	Medium	Medium	Medium	Mobile/attachment scanning critical [51]

Method (from taxonomy)	DMARC/SPF/DKIM	Secure email gateway & URL scanning	Identity telemetry (sign-ins, consent, token use)	Endpoint security (EDR)	User training/process controls	Notes
Smishing	N/A	Low	Medium	Low–Medium	Medium	Often outside corporate email controls [52]
Vishing / helpdesk SE	N/A	N/A	High	Medium	High	Best detection often comes from process + identity events [53]
Deepfake-enabled fraud	N/A	N/A	Low–Medium	Low	High	Detection is primarily procedural today [54]
Homograph/IDN attacks	Low	Medium	Medium	Medium	Medium	Browser display + domain controls help [55]
BEC / VEC	Medium	Medium	High	Low	High	Financial controls (verification, dual control) are core [2]
ClickFix (“copy/paste”)	N/A	Medium	Medium	High	High	EDR + script controls + training decisive [56]

Case studies and notable incidents 2019–2026

This section highlights representative incidents that illustrate *mechanisms*, not just headlines. Dates are included to clarify sequencing.

2019: deepfake voice fraud enters mainstream awareness

Business press reported a 2019 case in which criminals used AI-generated voice impersonation to convince a CEO to transfer roughly $243,000 (≈€220,000) in a CEO-fraud-style scenario. [36] While details vary across secondary reporting, the lasting lesson is that “voice = verification” became materially weaker after 2019, especially for high-trust payment workflows.

2020: Twitter[57] account hijacking via phone spear phishing of employees

On July 15, 2020, attackers used phone spear phishing/social engineering against employees to access internal tools and compromise prominent accounts. Twitter's own incident update describes a phone spear phishing attack targeting a small number of employees and requiring access to internal systems and specific credentials. [58] The New York State Department of Financial Services[59] investigation report provides additional detail, describing attackers calling employees while impersonating IT help desk and exploiting remote-work VPN troubleshooting expectations. [60] The United States Department of Justice[61] announced charges on July 31, 2020. [62]

Why this matters for phishing taxonomy: this is a vishing/helpdesk pretext that produced privileged access—an archetype now repeated in later helpdesk-targeting campaigns.

2021: cloud post-compromise detection responds to identity-centric attacks

CISA's April 15, 2021 advisory AA21-008A ("Detecting Post-Compromise Threat Activity in Microsoft Cloud Environments") includes the Sparrow tool for detecting compromised accounts and applications in Azure/M365 environments. [42] Although tied to a broader supply chain context, it demonstrates government recognition that **identity artifacts and cloud audit logs** are key battlegrounds.

2022: Twilio[63] and smishing-driven workforce credential theft

On August 8–9, 2022, Cloudflare described a targeted SMS phishing campaign similar to the one Twilio disclosed, aimed at employees and designed to steal credentials via fake login pages. [64] Reporting summarized that multiple Twilio employees were duped and attackers gained access to internal systems containing customer data. [65] The incident also had downstream ecosystem effects, with Signal[66] describing impacts related to Twilio's phishing attack on phone number verification services. [67]

Mechanism lessons: smishing works as a **device-surface and attention-surface exploit**; credential capture is often paired with SSO keywords (Okta/SSO) and "IT support" framing.

2022: Uber[68] and MFA fatigue + social engineering

Multiple analyses of the September 2022 Uber breach describe MFA "push bombing/fatigue" combined with social engineering to induce approval of a login attempt, enabling access to internal systems. [69] The persistent theme is that interactive, repeated prompts plus human deception can defeat MFA methods that are not phishing-resistant.

2023: MGM Resorts International[70] and Caesars Entertainment[71]—service desk and vendor-focused social engineering

Public reporting and filings indicate Caesars confirmed a social engineering attack starting in mid-August 2023 leading to data theft, and the broader incident cluster is widely linked to helpdesk or vendor-targeting social engineering. [72] These incidents underscore that organizational "support processes" function as authentication surfaces; attackers exploit them with high-context pretexts.

2024: deepfake video conference fraud at scale in Hong Kong[73]

Reuters (April 11, 2024) described a Hong Kong police-reported incident in which deepfake technology simulated a video conference, tricking an employee into transferring over $25 million. [74] Additional security vendor commentary framed it as a watershed moment for deception-driven fraud. [75]

Key insight: deepfakes amplify the credibility of "multi-party confirmation," collapsing a common anti-BEC control ("validate with the CFO and another colleague") if validation occurs within the attacker-controlled channel.

2024–2025: device code phishing shifts from niche to rapidly adopted

Microsoft reported an active device code phishing campaign by Storm-2372 (February 13, 2025) active since August 2024, using lures resembling messaging app experiences (WhatsApp, Signal, Microsoft Teams). [76] Microsoft's Digital Defense Report 2025 notes device code phishing's high risk and describes attackers prompting device code entry in Teams-style invitations, and also identifies nation-state and criminal adoption. [16]

2025: SaaS-focused vishing and connected-app abuse (UNC6040 / Salesforce)

Google's June 4, 2025 report on UNC6040 describes a financially motivated cluster specializing in vishing to compromise Salesforce instances, including tricking victims into authorizing a malicious connected app and sometimes requesting MFA codes during calls; extortion may follow months later. [21] A September 12, 2025 IC3 cyber advisory similarly describes UNC6040 leveraging vishing to access Salesforce accounts. [77]

This case bridges multiple taxonomy categories: vishing + OAuth-like connected app authorization + extortion monetization + potential actor collaboration.

2025: ClickFix and user-executed payload delivery

Microsoft's August 21, 2025 analysis documents ClickFix as a social engineering technique that tricks users into executing malicious commands, bypassing conventional automated security controls by inserting user interaction into the chain. [26] Government and industry alerts in 2024–2025 (including health-sector advisories and national alerts) indicate broad targeting patterns. [78]

2025: Scanception and QR-in-PDF credential harvesting

Broadcom describes Scanception (July 18, 2025) as QR codes embedded in PDF attachments redirecting users to fake Microsoft 365 login pages for credential harvesting, illustrating how attackers adapt to URL scanning controls. [24] ENISA's 2025 threat landscape also references QR code phishing used to evade email filtering. [79]

2026: OAuth redirect abuse, Tycoon2FA disruption, and AI-enabled phishing automation

- On March 2, 2026, Microsoft reported OAuth redirection abuse used in phishing to redirect victims to attacker infrastructure without token theft in that flow, targeting government and public-sector organizations. [22]
- On March 4, 2026, Microsoft and partners announced disruption of Tycoon2FA, described as active since at least 2023 and enabling MFA-defeating impersonation. [80] Independent reporting

from CrowdStrike notes Europol announced disruption and details domain seizures, illustrating law-enforcement/industry operations against phishing-as-a-service infrastructure. [81]

- On April 6, 2026, Microsoft described AI-enabled device code phishing with hyper-personalized lures and just-in-time code generation to defeat expiration windows. [19]

Threat actor profiles and motivations (pattern-based)

Threat actor ecosystems in modern phishing are increasingly **modular**:

- **State-aligned espionage** uses identity-flow abuse for stealthy access. Microsoft reports device code phishing observed among nation-state actors (Russia, Iran, China) as well as criminal groups. [16]
- **Financially motivated extortion and data theft** increasingly targets SaaS control planes and identity administrators. Google's UNC6040 reporting describes vishing-led Salesforce compromise and subsequent extortion behaviors. [21]
- **Phishing-as-a-service operators** productize kits (AiTM and device flow) and sell access at scale. Microsoft frames Tycoon2FA as an AiTM kit enabling broad actor use. [82]
- **Initial access brokers and handoff economics** shrink time-to-monetization and push attackers toward repeatable identity compromise. Mandiant's reporting on changing intrusion vectors and operational tempo contextualizes these economic pressures. [83]

Mitigation, incident response, and future trends

Defense against modern phishing is best framed as **reducing identity attack surface**, **hardening approval workflows**, and **building fast containment muscle** for token/session compromise—while recognizing that user training is necessary but not sufficient.

Technical controls (prioritized)

Phishing-resistant authentication (FIDO2/WebAuthn, PKI) as a structural control.
NIST SP 800-63B (Revision 4, 2025) strengthens requirements around phishing-resistant authentication options, including at higher assurance levels, and defines phishing resistance in authenticator requirements. [84] The U.S. government's phishing-resistant authenticator playbook (February 2024) explicitly categorizes OTPs (SMS/email/app), push notifications, and passwords as phishable, and identifies PKI and FIDO2 as leading phishing-resistant options. [85]

App governance: treat OAuth consents and connected apps as privileged events.
Because OAuth consent phishing can bypass MFA and persist beyond password resets, mitigation must include: - restricting user consent (admin approval workflow), - continuous review of granted permissions/scopes, - blocking or alerting on suspicious/new apps, - monitoring for token issuance anomalies and device code sign-ins. [86]

Device code flow controls.
Given observed criminal and state usage and the 2026 automation patterns, organizations should explicitly decide whether to: - disable or restrict device code flow where unnecessary, - scope device code flow to trusted device classes or networks, - alert on device code authentication events, especially in unusual contexts. The risk profile is supported by Microsoft and Proofpoint observations of device code phishing adoption. [87]

QR-code and mobile pathway hardening.
Adopt attachment sandboxing that can extract and detonate QR payloads, apply mobile DNS/web filtering, and enforce managed browsers on corporate mobile devices—because QR phish often moves the click to mobile surfaces. The growth of QR shortener traffic and QR-in-PDF campaigns supports prioritization. [88]

Defend helpdesks and recovery as high-risk authentication surfaces.
Modern vishing targets helpdesks for credential resets and MFA resets. Google and Microsoft both highlight helpdesk impersonation patterns driving compromise. [89] Defenses include: - verified callback directories (not caller ID), - step-up proof for resets (phishing-resistant factors), - time-delayed resets for privileged accounts, - “no MFA reset via phone” policies for high-risk roles.

Endpoint controls for “phishing that becomes execution.”
ClickFix exemplifies phishing leading to endpoint execution. Use PowerShell/script controls, EDR detections for suspicious clipboard-driven commands, and user education specifically about “copy-paste fixes.” [90]

Recommended controls mapping (practical)

| Phishing method | Prevent | Detect | Respond | |---|---|---| | AiTM / session theft | Phishing-resistant MFA; conditional access; session protection | Alerts on token replay patterns; impossible travel; new device + mailbox anomalies | Revoke sessions/tokens; reset auth methods; hunt mailbox rules | | OAuth device code phishing | Restrict device code flow; phishing-resistant MFA | Monitor device code sign-ins; correlate with email/chats | Revoke refresh tokens; investigate OAuth grants; lateral movement hunt | | OAuth consent / connected app | Disable user consent; app allowlists | Alert on new grants, high-privilege scopes | Remove app grants; rotate secrets; review service principals | | OAuth redirect abuse | Web isolation; attachment controls; endpoint hardening | Detect legit-login-to-malicious-redirect patterns | Contain endpoints; trace downloaded payloads; block redirect domains | | QR-code phishing | QR extraction + scanning; mobile web filtering | QR scanning telemetry; DNS logs | Credential reset if entered; token revoke; user outreach | | Vishing/helpdesk | Harden recovery; verified callback; separation of duties | Flag unusual reset volume; privileged reset requests | Freeze account changes; audit support tooling; post-incident training | | BEC/VEC | Payment verification and dual control; vendor change controls | Detect mailbox forwarding, unusual invoice patterns | Freeze payments; legal/finance coordination; notify bank quickly |

Incident response playbooks (identity-first)

A modern phishing IR playbook should assume “credential reset alone” may be insufficient because tokens and app grants persist. CISA’s Sparrow tooling focus (compromised accounts and applications) aligns with this identity-first posture. [42]

A concise identity-centric IR sequence (adapt as needed):

1. **Contain identity sessions**: revoke refresh tokens and active sessions; disable suspicious accounts; enforce reauthentication.
2. **Audit OAuth grants and connected apps**: remove unknown apps; review scopes; rotate app secrets/certificates.
3. **Hunt mailbox rule abuse**: forwarding rules, hidden inbox rules, OAuth mail access, and unusual eDiscovery/export.

4. **Scope lateral movement**: internal phishing, new admin consent events, privileged role assignment changes.
5. **Endpoint validation** (where phishing delivered payload): isolate affected hosts; examine downloads and persistence.
6. **Finance workflow containment** for BEC: hold payments, validate vendor banking changes using independent channels.

Legal and regulatory considerations (high-level, not legal advice)

Phishing often triggers regulatory obligations when it results in system compromise, service disruption, or personal data exposure.

- **GDPR** requires notifying the supervisory authority of a personal data breach "where feasible, not later than 72 hours" after becoming aware, unless unlikely to result in risk to rights and freedoms. [91]
- **NIS2 (Directive (EU) 2022/2555)** introduces accelerated incident reporting expectations (including early warning "within 24 hours" for significant incidents) as part of broader risk-management obligations. [92]
- **DORA** applies to EU financial entities and becomes applicable from January 17, 2025, strengthening ICT risk and incident reporting regimes in the financial sector. [93]
- The **U.S. SEC** cybersecurity disclosure rules (adopted July 26, 2023) require Form 8-K Item 1.05 disclosure of material cybersecurity incidents generally within four business days after materiality determination. [94]

Operational implication: incident response for phishing must be integrated with legal/GRC early, because token theft and mailbox compromise can create uncertainty about scope and data exposure timelines.

Future trends and research gaps

Shift to interactive and multimodal social engineering.
Mandiant's reporting suggests a continued decline in email phishing's share (in its incident sets) with increases in voice-based methods. [95] Google's 2026 cloud threat horizons similarly calls out vishing shares and specific financially motivated clusters. [3]

Agentic AI and automation of the full phishing lifecycle.
Recent peer-reviewed discussion argues agentic AI can enable autonomous planning, multi-channel execution, and adaptive feedback loops—expanding attack scalability and adaptability. [96] Microsoft's April 2026 reporting provides a real-world anchor: automation targeted not only content but timing and infrastructure to keep device codes valid. [19]

Deepfake detection realism as a core gap.
Deepfake voice detection research highlights that real-world performance depends heavily on modeling the presentation channel; dataset realism and evaluation methodology are key research bottlenecks. [37]

Cross-surface telemetry correlation remains under-validated.
Many identity-centric attacks require correlating email, IdP logs, SaaS audit logs, mobile telemetry, and

endpoint signals. Tooling exists (e.g., CISA Sparrow for Azure/M365), but operational maturity and standardized detection engineering patterns remain uneven. [97]

Actionable recommendations and concise defender checklist

Actionable recommendations (condensed into high-impact moves):

- Implement phishing-resistant authentication (FIDO2/WebAuthn and/or PKI) for privileged users and financial workflows, using phased rollout patterns consistent with federal playbooks and NIST's emphasis on phishing resistance. [98]
- Treat OAuth grants, device code sign-ins, and connected-app authorizations as privileged security events; restrict user consents and continuously audit app permissions. [99]
- Build QR visibility: extract QR payloads from attachments; extend web protections to mobile; monitor QR shortener usage and block suspicious QR services as needed. [46]
- Harden helpdesks and recovery: implement verified callback procedures, prohibit phone-based MFA resets for high-risk accounts, and log/alert on reset anomalies. [100]
- Update training content to modern patterns: device code flow abuse, AiTM "login succeeds but attacker wins," QR-in-PDF, and ClickFix copy-paste execution. [101]
- Prepare identity-first IR runbooks: token revocation, app consent review, mailbox rule hunting, and SaaS audit scoping—beyond password resets. [102]

Concise checklist (operational):

- [ ] Phishing-resistant MFA enabled for admins, finance, execs; legacy MFA minimized for high-risk workflows. [98]
- [ ] User OAuth consent restricted; admin consent workflow enforced; app grants continuously reviewed. [99]
- [ ] Device code flow policy decided (allowed only where needed) and monitored. [48]
- [ ] Mailbox forwarding rules and OAuth mail scopes monitored with alerts and weekly audits. [97]
- [ ] QR extraction/scanning on attachments; mobile web protections deployed. [51]
- [ ] Helpdesk reset procedures hardened (verified callback, step-up identity proof, privileged reset controls). [103]
- [ ] ClickFix defenses: restrict script execution, monitor "paste and run," and train users to treat copy/paste "fixes" as suspicious. [104]
- [ ] IR playbooks include token revocation + OAuth grant review + endpoint isolation for redirect/malware delivery. [105]
- [ ] Regulatory notification pathways pre-planned (GDPR/NIS2/DORA/SEC as applicable). [106]

References

- Anti-Phishing Working Group[8]. *Phishing Activity Trends Report, 1st Quarter 2025* (PDF). July 2, 2025. [23]
- Anti-Phishing Working Group[8]. *Phishing Activity Trends Report, 4th Quarter 2025* (PDF). February 18, 2026. [107]
- Cybersecurity and Infrastructure Security Agency[41]. *Detecting Post-Compromise Threat Activity in Microsoft Cloud Environments (AA21-008A).* April 15, 2021. [108]

- Cybersecurity and Infrastructure Security Agency[41]. *Sparrow (GitHub repository).* (Accessed via GitHub). [109]
- European Union Agency for Cybersecurity[110]. *ENISA Threat Landscape 2025* (PDF). October 2025. [79]
- National Institute of Standards and Technology[10]. *Trustworthy Email (SP 800-177)* (PDF). October 2016. [111]
- National Institute of Standards and Technology[10]. Temoshok et al. *Digital Identity Guidelines: Authentication and Authenticator Management (SP 800-63B, Rev. 4)* (PDF). 2025. [112]
- Internet Engineering Task Force[113]. *RFC 8628: OAuth 2.0 Device Authorization Grant.* August 2019. [15]
- Internet Engineering Task Force[113]. *RFC 6749: The OAuth 2.0 Authorization Framework.* October 2012. [114]
- Internet Engineering Task Force[113]. *RFC 7489: Domain-based Message Authentication, Reporting, and Conformance (DMARC).* March 2015. [115]
- Internet Engineering Task Force[113]. *RFC 6376: DomainKeys Identified Mail (DKIM) Signatures.* September 2011. [116]
- Internet Engineering Task Force[113]. *RFC 7208: Sender Policy Framework (SPF).* April 2014. [117]
- Internet Engineering Task Force[113]. *RFC 5890: Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework.* August 2010. [118]
- Unicode Consortium[119]. *Unicode Technical Report #36: Unicode Security Considerations.* (Referenced for spoofing/confusables discussion). [44]
- Microsoft. *Microsoft Digital Defense Report 2025: Safeguarding Trust in the AI Era* (PDF). 2025. [120]
- Microsoft. *Storm-2372 conducts device code phishing campaign.* Microsoft Security Blog, February 13, 2025. [121]
- Microsoft. *Defending against evolving identity attack techniques.* Microsoft Security Blog, May 29, 2025. [122]
- Microsoft. *Disrupting threats targeting Microsoft Teams.* Microsoft Security Blog, October 7, 2025. [123]
- Microsoft. *Think before you Click(Fix): Analyzing the ClickFix social engineering technique.* Microsoft Security Blog, August 21, 2025. [26]
- Microsoft. *OAuth redirection abuse enables phishing and malware delivery.* Microsoft Security Blog, March 2, 2026. [22]
- Microsoft. *Inside Tycoon2FA: How a leading AiTM phishing kit operated at scale.* Microsoft Security Blog, March 4, 2026. [124]
- Microsoft. *How a global coalition disrupted Tycoon 2FA.* Microsoft On the Issues, March 4, 2026. [80]
- Microsoft. *Inside an AI-enabled device code phishing campaign.* Microsoft Security Blog, April 6, 2026. [125]
- Mandiant[126]. *M-Trends 2025* (PDF). 2025. [127]
- Mandiant[126] (via Google Cloud). *M-Trends 2026: Data, Insights, and Strategies From the Frontlines.* March 23, 2026. [39]

- Google Cloud / GTIG. *The Cost of a Call: From Voice Phishing to Data Extortion*. June 4, 2025. [21]
- Google Cloud. *Cloud Threat Horizons Report H1 2026*. (Web resource). [3]
- Palo Alto Networks[128] (Unit 42). *Phishing on the Edge of the Web and Mobile Using QR Codes*. (Web article). [25]
- Broadcom[129]. *Scanception: A sophisticated QR Code phishing campaign*. July 18, 2025. [130]
- Proofpoint. *Access Granted: Phishing with device code authorization for account takeover*. December 18, 2025. [18]
- Proofpoint. *Disruption targets Tycoon 2FA, popular AiTM PhaaS*. March 4, 2026. [131]
- Cloudflare[132]. *The mechanics of a sophisticated phishing scam and how we stopped it*. August 9, 2022. [64]
- European Union[133]. *Directive (EU) 2022/2555 (NIS2)* (EUR-Lex). December 27, 2022. [92]
- European Securities and Markets Authority[134]. *Digital Operational Resilience Act (DORA) overview*. (Web page). [135]
- United States Securities and Exchange Commission[136]. *SEC adopts rules on cybersecurity risk management, strategy, governance, and incident disclosure*. July 26, 2023. [137]
- Reuters. *Real insurance coverage for increasing AI deepfake risks* (includes Hong Kong deepfake video call case). April 11, 2024. [74]
- New York State Department of Financial Services[59]. *Twitter Investigation Report*. October 14, 2020. [138]
- Twitter[57]. *An update on our security incident*. July 30, 2020. [58]
- Delgado et al. *On Deepfake Voice Detection — It's All in the Presentation* (arXiv PDF). 2025. [139]
- Wang et al. *Detecting camouflaged IDN-based phishing attacks via Siamese neural network* (ScienceDirect abstract page). 2024. [45]
- Yao et al. *The Psychological Manipulation of Phishing Emails: A Cognitive Bias Approach* (open access). 2025. [140]
- Khadka et al. *A Survey on the Principles of Persuasion as a Social Engineering Approach* (PDF). 2024. [30]

---

[1] [16] [17] [20] [48] [49] [66] [68] [71] [86] [87] [99] [120] [126] https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Microsoft-Digital-Defense-Report-2025.pdf

https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Microsoft-Digital-Defense-Report-2025.pdf

[2] [39] [83] [95] [132] https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026

https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026

[3] https://cloud.google.com/security/report/resources/cloud-threat-horizons-report-h1-2026

https://cloud.google.com/security/report/resources/cloud-threat-horizons-report-h1-2026

[4] [14] [82] [124] https://www.microsoft.com/en-us/security/blog/2026/03/04/inside-tycoon2fa-how-a-leading-aitm-phishing-kit-operated-at-scale/

https://www.microsoft.com/en-us/security/blog/2026/03/04/inside-tycoon2fa-how-a-leading-aitm-phishing-kit-operated-at-scale/

[5] [15] https://datatracker.ietf.org/doc/html/rfc8628

https://datatracker.ietf.org/doc/html/rfc8628

[6] [24] [38] [51] [63] [130] https://www.broadcom.com/support/security-center/protection-bulletin/scanception-a-sophisticated-qr-code-phishing-campaign

https://www.broadcom.com/support/security-center/protection-bulletin/scanception-a-sophisticated-qr-code-phishing-campaign

[7] [23] [33] https://docs.apwg.org/reports/apwg_trends_report_q1_2025.pdf

https://docs.apwg.org/reports/apwg_trends_report_q1_2025.pdf

[8] [81] https://www.crowdstrike.com/en-us/blog/tycoon2fa-phishing-as-a-service-platform-persists-following-takedown/

https://www.crowdstrike.com/en-us/blog/tycoon2fa-phishing-as-a-service-platform-persists-following-takedown/

[9] [107] https://docs.apwg.org/reports/apwg_trends_report_q4_2025.pdf

https://docs.apwg.org/reports/apwg_trends_report_q4_2025.pdf

[10] [58] [136] https://blog.x.com/en_us/topics/company/2020/an-update-on-our-security-incident

https://blog.x.com/en_us/topics/company/2020/an-update-on-our-security-incident

[11] [13] [111] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-177.pdf?utm_source=Securitylabru

https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-177.pdf?utm_source=Securitylabru

[12] [117] https://datatracker.ietf.org/doc/html/rfc7208

https://datatracker.ietf.org/doc/html/rfc7208

[18] https://www.proofpoint.com/us/blog/threat-insight/access-granted-phishing-device-code-authorization-account-takeover

https://www.proofpoint.com/us/blog/threat-insight/access-granted-phishing-device-code-authorization-account-takeover

[19] [28] [125] https://www.microsoft.com/en-us/security/blog/2026/04/06/ai-enabled-device-code-phishing-campaign-april-2026/

https://www.microsoft.com/en-us/security/blog/2026/04/06/ai-enabled-device-code-phishing-campaign-april-2026/

[21] [53] [73] [89] https://cloud.google.com/blog/topics/threat-intelligence/voice-phishing-data-extortion

https://cloud.google.com/blog/topics/threat-intelligence/voice-phishing-data-extortion

[22] [50] [105] [129] https://www.microsoft.com/en-us/security/blog/2026/03/02/oauth-redirection-abuse-enables-phishing-malware-delivery/

https://www.microsoft.com/en-us/security/blog/2026/03/02/oauth-redirection-abuse-enables-phishing-malware-delivery/

[25] [46] [88] https://unit42.paloaltonetworks.com/qr-codes-as-attack-vector/

https://unit42.paloaltonetworks.com/qr-codes-as-attack-vector/

[26] [56] [90] [101] [104] https://www.microsoft.com/en-us/security/blog/2025/08/21/think-before-you-clickfix-analyzing-the-clickfix-social-engineering-technique/

https://www.microsoft.com/en-us/security/blog/2025/08/21/think-before-you-clickfix-analyzing-the-clickfix-social-engineering-technique/

[27] https://attack.mitre.org/techniques/T1204/004/

https://attack.mitre.org/techniques/T1204/004/

[29] [140] https://www.sciencedirect.com/org/science/article/pii/S1546221825009968

https://www.sciencedirect.com/org/science/article/pii/S1546221825009968

[30] [134] https://arxiv.org/pdf/2412.18488

https://arxiv.org/pdf/2412.18488

[31] https://www.emerald.com/ics/article/33/5/845/1267929/Persuasion-under-pressure-the-influence-of

https://www.emerald.com/ics/article/33/5/845/1267929/Persuasion-under-pressure-the-influence-of

[32] [76] [121] https://www.microsoft.com/en-us/security/blog/2025/02/13/storm-2372-conducts-device-code-phishing-campaign/

https://www.microsoft.com/en-us/security/blog/2025/02/13/storm-2372-conducts-device-code-phishing-campaign/

[34] https://arxiv.org/html/2508.21457v1

https://arxiv.org/html/2508.21457v1

[35] [41] [54] [74] https://www.reuters.com/legal/legalindustry/real-insurance-coverage-increasing-ai-deepfake-risks-2024-04-11/

https://www.reuters.com/legal/legalindustry/real-insurance-coverage-increasing-ai-deepfake-risks-2024-04-11/

[36] [113] https://www.forbes.com/sites/jessedamiani/2019/09/03/a-voice-deepfake-was-used-to-scam-a-ceo-out-of-243000/

https://www.forbes.com/sites/jessedamiani/2019/09/03/a-voice-deepfake-was-used-to-scam-a-ceo-out-of-243000/

[37] [139] https://arxiv.org/pdf/2509.26471

https://arxiv.org/pdf/2509.26471

[40] [57] [79] https://www.enisa.europa.eu/sites/default/files/2025-11/ENISA%20Threat%20Landscape%202025.pdf

https://www.enisa.europa.eu/sites/default/files/2025-11/ENISA%20Threat%20Landscape%202025.pdf

[42] [59] [102] [108] https://www.cisa.gov/news-events/cybersecurity-advisories/aa21-008a

https://www.cisa.gov/news-events/cybersecurity-advisories/aa21-008a

[43] [55] [118] https://datatracker.ietf.org/doc/html/rfc5890

https://datatracker.ietf.org/doc/html/rfc5890

[44] https://www.unicode.org/reports/tr36/tr36-6.html

https://www.unicode.org/reports/tr36/tr36-6.html

[45] https://www.sciencedirect.com/science/article/abs/pii/S0167404823005783

https://www.sciencedirect.com/science/article/abs/pii/S0167404823005783

[47] https://www.researchgate.net/publication/395022264_Automatically_Detecting_Voice_Phishing_A_Large_Audio_Model_Approach

https://www.researchgate.net/publication/395022264_Automatically_Detecting_Voice_Phishing_A_Large_Audio_Model_Approach

[52] [64] https://blog.cloudflare.com/2022-07-sms-phishing-attacks/

https://blog.cloudflare.com/2022-07-sms-phishing-attacks/

[60] [100] [103] [138] https://www.dfs.ny.gov/Twitter_Report

https://www.dfs.ny.gov/Twitter_Report

[61] [91] [106] https://gdpr-info.eu/art-33-gdpr/

https://gdpr-info.eu/art-33-gdpr/

[62] https://www.justice.gov/archives/opa/pr/three-individuals-charged-alleged-roles-twitter-hack

https://www.justice.gov/archives/opa/pr/three-individuals-charged-alleged-roles-twitter-hack

[65] https://www.cybersecuritydive.com/news/twilio-phishing-attack/629142/

https://www.cybersecuritydive.com/news/twilio-phishing-attack/629142/

[67] https://support.signal.org/hc/en-us/articles/4850133017242-Twilio-Incident-What-Signal-Users-Need-to-Know

https://support.signal.org/hc/en-us/articles/4850133017242-Twilio-Incident-What-Signal-Users-Need-to-Know

[69] https://assets.kpmg.com/content/dam/kpmgsites/in/pdf/2022/09/27-september-2022-lessons-to-learn-from-the-uber-security-breach.pdf.coredownload.inline.pdf

https://assets.kpmg.com/content/dam/kpmgsites/in/pdf/2022/09/27-september-2022-lessons-to-learn-from-the-uber-security-breach.pdf.coredownload.inline.pdf

[70] [80] [133] https://blogs.microsoft.com/on-the-issues/2026/03/04/how-a-global-coalition-disrupted-tycoon/

https://blogs.microsoft.com/on-the-issues/2026/03/04/how-a-global-coalition-disrupted-tycoon/

[72] https://www.cybersecuritydive.com/news/caesars-social-engineering-breach/695995/

https://www.cybersecuritydive.com/news/caesars-social-engineering-breach/695995/

[75] https://www.trendmicro.com/en_us/research/24/b/deepfake-video-calls.html

https://www.trendmicro.com/en_us/research/24/b/deepfake-video-calls.html

[77] https://www.ic3.gov/CSA/2025/250912.pdf

https://www.ic3.gov/CSA/2025/250912.pdf

[78] https://www.hhs.gov/sites/default/files/clickfix-attacks-sector-alert-tlpclear.pdf

https://www.hhs.gov/sites/default/files/clickfix-attacks-sector-alert-tlpclear.pdf

[84] [112] [119] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63B-4.pdf

https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63B-4.pdf

[85] [98] https://www.idmanagement.gov/playbooks/altauthn/

https://www.idmanagement.gov/playbooks/altauthn/

[92] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX%3A32022L2555

https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX%3A32022L2555

[93] [135] https://www.esma.europa.eu/esmas-activities/digital-finance-and-innovation/digital-operational-resilience-act-dora

https://www.esma.europa.eu/esmas-activities/digital-finance-and-innovation/digital-operational-resilience-act-dora

[94] [128] [137] https://www.sec.gov/newsroom/press-releases/2023-139

https://www.sec.gov/newsroom/press-releases/2023-139

[96] https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2026.1795045/full

https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2026.1795045/full

[97] [109] [110] https://github.com/cisagov/Sparrow

https://github.com/cisagov/Sparrow

[114] https://datatracker.ietf.org/doc/html/rfc6749

https://datatracker.ietf.org/doc/html/rfc6749

[115] https://datatracker.ietf.org/doc/html/rfc7489

https://datatracker.ietf.org/doc/html/rfc7489

[116] https://datatracker.ietf.org/doc/html/rfc6376

https://datatracker.ietf.org/doc/html/rfc6376

[122] https://www.microsoft.com/en-us/security/blog/2025/05/29/defending-against-evolving-identity-attack-techniques/

https://www.microsoft.com/en-us/security/blog/2025/05/29/defending-against-evolving-identity-attack-techniques/

[123] https://www.microsoft.com/en-us/security/blog/2025/10/07/disrupting-threats-targeting-microsoft-teams/

https://www.microsoft.com/en-us/security/blog/2025/10/07/disrupting-threats-targeting-microsoft-teams/

[127] https://services.google.com/fh/files/misc/m-trends-2025-en.pdf

https://services.google.com/fh/files/misc/m-trends-2025-en.pdf

[131] https://www.proofpoint.com/us/blog/threat-insight/disruption-targets-tycoon-2fa-popular-aitm-phaas

https://www.proofpoint.com/us/blog/threat-insight/disruption-targets-tycoon-2fa-popular-aitm-phaas